

SPECTRAL GAP IN THE GROUP OF AFFINE TRANSFORMATIONS OVER PRIME FIELDS

ELON LINDENSTRAUSS AND PÉTER P. VARJÚ

ABSTRACT. We study random walks on the groups $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. We estimate the spectral gap in terms of the spectral gap of the projection to the linear part $\mathrm{SL}_d(\mathbb{F}_p)$. This problem is motivated by an analogue in the group $\mathbb{R}^d \rtimes \mathrm{SO}(d)$, which have application to smoothness of self-similar measures.

1. INTRODUCTION

Let G be a finite group. Fix a set $S \subset G$, and let $X_1, X_2 \dots \in S$ be a sequence of independent random elements taking each element of S with equal probability. Denote the product of the first l by $Y_l = X_l \cdots X_1$. The sequence Y_1, \dots, Y_l is called the (simple) random walk on G generated by S .

We consider the following operator acting on the space $\mathbb{C}[G]$ of complex valued functions on G :

$$\mathcal{L}(G, S)f(g) = \frac{1}{|S|} \sum_{s \in S} f(s^{-1}g),$$

for $f \in \mathbb{C}[G]$ and $g \in G$. In addition, we consider $\mathcal{L}_0(G, S)$, the restriction of $\mathcal{L}(G, S)$ to the one codimensional subspace of $\mathbb{C}[G]$ consisting of functions orthogonal to the constants. This averaging operator is intimately connected with the random walk Y_l , and in particular the norm of $\mathcal{L}_0(G, S)$ is closely connected with how quickly this random walk becomes equidistributed. Clearly $\|\mathcal{L}_0(G, S)\| \leq \|\mathcal{L}(G, S)\| = 1$. We shall call the difference $1 - \|\mathcal{L}_0(G, S)\|$ the *spectral gap of the random walk*. If S is symmetric the spectral gap coincides with difference

Date: September 7, 2016.

Key words and phrases. spectral gap, random walk, random walks on groups, expanders.

EL was supported by the European Research Council (Advanced Research Grant 267259) and the ISF (grant 983/09) PV was supported by the Simons Foundation and the European Research Council (Advanced Research Grant 267259). The authors would like to thank the Israeli Institute for Advanced Study for its hospitality during the fall of 2013.

between the trivial eigenvalue 1 of $\mathcal{L}(G, S)$ and the greatest eigenvalue of the operator $\|\mathcal{L}_0(G, S)\|$, though in general (despite the name which seems to be fairly standard) what we call the spectral gap has no direct spectral interpretation. The operator norm here and everywhere below is with respect to the L^2 norm on the space the operator is acting on, in this case the finite group G equipped with the counting measure.

It is easily seen that the random walk mixes rapidly if the spectral gap is large. Indeed, denote by $\delta_1 \in \mathbb{C}[G]$, the function given by $\delta_1(1) = 1$ and $\delta_1(g) = 0$ for $g \neq 1$, where 1 denotes the multiplicative unit in G and in \mathbb{C} and in any multiplicative group. Then one can show by induction, that for all integer $l \geq 0$, the probability that $Y_l = g$ is $\mathcal{L}(G, S)^l \delta_1(g)$. We can write

$$\delta_1(g) = \frac{1}{|G|} + f(g),$$

where f is a function orthogonal to the constants. Then

$$\left\| \frac{1}{|G|} - \mathcal{L}(S)^l \delta_1 \right\|_{L^\infty} \leq \left\| \frac{1}{|G|} - \mathcal{L}(S)^l \delta_1 \right\|_{L^2} \leq e^{-l(1-\|\mathcal{L}_0(S)\|)}.$$

In particular, the distribution of Y_l is very close to uniform if say $l \geq 10(1 - \|\mathcal{L}_0(S)\|)^{-1} \log |G|$. More precisely, for such an l , we have

$$\left| \mathbb{P}(Y_l = g) - \frac{1}{|G|} \right| \leq \frac{1}{|G|^{10}}.$$

There is also a combinatorial way to characterize large spectral gap. If the spectral gap is large, then the Cayley graph of G with respect to S has a large isoperimetric constant. If S is symmetric (i.e. $s \in S$ implies $s^{-1} \in S$), then the converse is also true. Graphs with large isoperimetric constants are called expanders. For more details we refer to Lubotzky's survey [Lub12].

The problem of studying spectral gaps of random walks is interesting in its own right, and it has been studied extensively. Recently, spectral gap estimates were used together with sieve techniques to prove various results in number theory and group theory. A detailed account on these developments would go beyond the scope of our paper, so we refer the interested reader to the recent surveys [Lub12] and [Kow12].

1.1. Statement of the result. Let p be a prime and denote by \mathbb{F}_p the finite field of order p . Our result compares the spectral gap of a random walk on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$, the group of affine transformations of \mathbb{F}_p^d with its projection to $\mathrm{SL}_d(\mathbb{F}_p)$.

Theorem 1. *There is a number c depending only on d such that the following holds. Let $S' \subset \mathrm{SL}_d(\mathbb{F}_p)$, and let $S \subset \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ be such a set that for each $g \in S$ there is precisely one $\sigma \in S'$ such that the linear part of g is σ . Suppose further that S is not contained in a coset of a proper subgroup of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. Then*

$$1 - \|\mathcal{L}_0(\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p), S)\| \geq c \cdot \min\{1 - \|\mathcal{L}_0(\mathrm{SL}_d(\mathbb{F}_p), S')\|, |S|^{-1}\}.$$

Up to the constant c , the bound is sharp, as can be seen by the example when all but one element of S is contained in a subgroup isomorphic to $\mathrm{SL}_d(\mathbb{F}_p)$. However, when the distribution of S is better among the cosets of such subgroups the bound can be improved. In the next section, we will formulate a slightly more general version of this theorem with an improved bound.

In order to apply Theorem 1, a bound on the spectral gap for the projection of the random walk in $\mathrm{SL}_d(\mathbb{F}_p)$ is required. The following important result of Bourgain and Gamburd provides such a bound:

Theorem A (Bourgain, Gamburd). *Let $\overline{S} \subset \mathrm{SL}_d(\mathbb{Z})$ be a finite symmetric set, which generates a Zariski-dense subgroup. Then there is a number c depending on \overline{S} (but not on p) such that the following holds for all but finitely many primes p . Let S be the mod p projection of \overline{S} . Then*

$$1 - \|\mathcal{L}_0(\mathrm{SL}_d(\mathbb{F}_p), S)\| > c.$$

The $d = 2$ case of this theorem is [BG08, Theorem 1], and this has been worked out by Kowalski [Kow13] with explicit constants. A key ingredient in the proof is Helfgott's product theorem in [Hel08]. The $d \geq 3$ case is [BG09, Theorem 1.2] which assumes a generalization of Helfgott's theorem as a black box. This generalization is due to Helfgott [Hel11] in the $d = 3$ case, and to Breuillard, Green and Tao [BGT11] and independently by Pyber and Szabó [PS10] in the general case.

A key problem in the subject highlighted in [LW93] is determining how the size of the spectral gap depends on the choice of generators. It seems plausible that the conclusion of Theorem A could hold with a constant c depending only on d and $|S|$, and not on a previously fixed set in $\mathrm{SL}_d(\mathbb{Z})$. The following theorem by Breuillard and Gamburd [BG10, Theorem 1.1] gives some evidence in this direction:

Theorem B (Breuillard, Gamburd). *For any $\delta > 0$ and positive integer N , there is a constant $c > 0$ depending only on δ and N such that for any sufficiently large X , for all but X^δ primes $p \leq X$, for any*

symmetric generating set $S \subset \mathrm{SL}_2(\mathbb{F}_p)$ with $|S| = N$,

$$1 - \|\mathcal{L}(\mathrm{SL}_2(\mathbb{F}_p), S)\| > c.$$

Theorem 1 above can be seen in this general context: while the spectral gap on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ is dependent on the choice of generators S' for $\mathrm{SL}_d(\mathbb{F}_p)$, the estimate given by the theorem is uniform in the way S' is lifted to a generating set S on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. If one takes S to be the projection mod p of a fixed set $\overline{S} \subset \mathbb{Z}^d \rtimes \mathrm{SL}_d(\mathbb{Z}^d)$ generating a (fixed) Zariski dense subgroup, establishing a spectral gap (uniform in p) for the corresponding averaging operator can be obtained by an adaptation of the method of Bourgain and Gamburd without introducing any substantial new ideas. In particular, it is a very special case of the main result of [SGV12, Theorem 1].

1.2. Motivation. One source of interest in the group $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ stems from a continuous analogue of the problem. In that analogue, the role of $\mathrm{SL}_d(\mathbb{F}_p)$ is played by the compact Lie group $\mathrm{SO}(d)$ and $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ is replaced by $\mathbb{R}^d \rtimes \mathrm{SO}(d)$, the group of orientation preserving isometries of Euclidean space. In the paper [LV14b], we prove an analogue of Theorem 1 in that setting. This has two applications of independent interest in quite different directions:

- Under the assumption that a corresponding random walk on $\mathrm{SO}(d)$ has spectral gap, we show that a self-similar measure is absolutely continuous, provided the contraction coefficients of the self-similarities are sufficiently close to 1.
- In [Var] a local-central limit theorem for a random walk on \mathbb{R}^d by Euclidean isometries is proved. In [LV14b] we strengthen this result when the underlying random walk on $\mathrm{SO}(d)$ has spectral gap to show that this local-central limit theorem holds at a scale exponentially small in the number of steps (when $d \geq 3$ and the rotation part of the random walk generates a dense subgroup of $\mathrm{SO}(d)$, the local-central limit theorem is established in [Var] only up to the scale $e^{-O(l^{1/3})}$ where l is the number of steps).

1.3. Ideas in the proof. We heavily exploit the method of Bourgain and Gamburd in our proof of Theorem 1. Note however that the product theorems of [Hel08, Hel11, BGT11, PS10] are not used in the proof of Theorem 1, at least not directly; in their stead we use the assumed spectral gap of the averaging operator $\mathcal{L}_0(\mathrm{SL}_d(\mathbb{F}_p), S')$ corresponding to the associated random walk on $\mathrm{SL}_d(\mathbb{F}_p)$.

We recall the essence of the method in Theorem D in Section 4. To apply this theorem to the problem at hand, we need to show that the

random walk does not concentrate on cosets of subgroups isomorphic to $\mathrm{SL}_d(\mathbb{F}_p)$. More precisely, we show that

$$\mathbb{P}(Y_l \in A) \leq 4p^{-d/4}$$

if A is a coset of a subgroup isomorphic to $\mathrm{SL}_d(\mathbb{F}_p)$ and $l \geq C \log p$ with a constant C sufficiently large.

This non-concentration estimate proved in Section 3 is the main new contribution in our paper. It is essentially equivalent to proving a non-concentration estimate for the random walk on \mathbb{F}_p^d generated by S , which we do by showing the L^2 -norm of the probability measure on \mathbb{F}_p^d after $O_{d,S'}(\log p)$ many steps becomes small.

The proof of this fact rests upon an observation that if η is an (arbitrary) probability measure on \mathbb{F}_p^d and the absolute value of its Fourier transform is almost constant in the appropriate sense on $\mathbb{F}_p^d \setminus \{0\}$, then the measure is either spread out on \mathbb{F}_p^d or most of the contribution to the L^2 -norm of η comes from a single atom; cf. Proposition 4.

Using this observation we argue iteratively: if the random walk on \mathbb{F}_p^d after some steps concentrates the L^2 -norm in a single atom, we can use that not all elements of S move this atom to the same point to show that the next step quantifiably reduces the L^2 -norm. If the Fourier transform does not have almost constant absolute value for all nonzero coefficients, we can use the spectral gap for the projection to $\mathrm{SL}_d(\mathbb{F}_p^d)$, to prove that the next step quantifiably reduces the L^4 norm of the Fourier transform.

To carry out this argument, we have to work with both L^2 and L^4 norms. Therefore it will be necessary, to relate the L^2 and L^4 spectral gaps. This can be done using either Riesz-Thorin interpolation as we do here, or some classical results on L^q spaces similarly to the paper [BFGM07] and the uniform convexity of L^p spaces [BL00, Chapter 9] as was done in an earlier version of this paper [LV14a].

The paper contains another new idea in Section 4.1, where we prove a result about growth of product-sets in the group $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. The result (Proposition 8) itself is not new; similar statements appear in the papers [SGV12] and [PS10]. However, we give a new proof that can be adapted to work in the continuous case as done in [LV14b].

1.4. An open problem. An interesting analogue of the results we obtained here is provided by a random walk using a set S of generators on the group $\mathrm{SL}_2(\mathbb{F}_p) \times \mathrm{SL}_2(\mathbb{F}_p)$. Is it possible to estimate the spectral gap in terms of the spectral gaps of the projections to the direct factors analogously to Theorem 1?

If one tries to prove such an estimate using the method of Bourgain and Gamburd then the following problem arises. The group $\mathrm{SL}_2(\mathbb{F}_p) \times \mathrm{SL}_2(\mathbb{F}_p)$ contains the subgroup $\{(g, g) : g \in \mathrm{SL}_2(\mathbb{F}_p)\}$ and its conjugates, and it may happen that the random walk concentrates too much mass on such a subgroup. If this obstacle could be ruled out, then a positive solution to the above problem would follow immediately from the method of Bourgain and Gamburd. This difficulty is similar to the one we tackle in this paper, but its solution probably require a different set of ideas. We also mention that when one looks at the problem in the group $\mathrm{SL}_2(\mathbb{F}_{p_1}) \times \mathrm{SL}_2(\mathbb{F}_{p_2})$ for different primes $p_1 \neq p_2$, then the problem disappears, since all proper subgroups of $\mathrm{SL}_2(\mathbb{F}_{p_1}) \times \mathrm{SL}_2(\mathbb{F}_{p_2})$ projects into a proper subgroup of one of the factors.

1.5. Organization. In the next section, we introduce some more notation and state a more technical and somewhat stronger version of Theorem 1. In Section 3, we prove the crucial non-concentration estimate mentioned above. We recall the method of Bourgain and Gamburd in Section 4 and use it to deduce our results.

Acknowledgment. We are grateful to the referee and Lam Pham for carefully reading the paper and for suggestions that significantly improved the presentation of the paper.

2. NOTATION

For $(v_1, \theta_1), (v_2, \theta_2) \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$, the product is defined by

$$(v_1, \theta_1) \cdot (v_2, \theta_2) = (v_1 + \theta_1 v_2, \theta_1 \cdot \theta_2).$$

If $g \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ can be written in the form $g = (v, \theta)$ then we write $v(g) = v$ and $\theta(g) = \theta$. In other words $v(g)$ and $\theta(g)$ are the projections to the factors \mathbb{F}_p^d and $\mathrm{SL}_d(\mathbb{F}_p)$ respectively. We note that $v(g)$ is not intrinsically defined and we fix one choice for the entire paper.

The group $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ naturally acts on \mathbb{F}_p^d by means of the formula $g.x = v(g) + \theta(g)x$ for $g \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ and $x \in \mathbb{F}_p^d$. This action is consistent with the above product law, i.e. we have the identity $(g_1 \cdot g_2).x = g_1.(g_2.x)$.

It is easy to check that the inverse of an element $g \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ is given by the formula

$$(1) \quad g^{-1} = (-\theta(g)^{-1}.v(g), \theta(g)^{-1}).$$

We identify measures on finite sets with their Radon-Nikodym derivative with respect to the counting measure. Thus the difference between our use of the words function and measure is purely rhetoric. With this convention we also write $f(A) = \sum_{x \in A} f(x)$, where A is

a finite set and f is a function (or measure) defined on a finite set containing A . A probability measure is a non-negative measure with total mass 1. The Dirac delta measure concentrated at the point x is a probability measure δ_x such that $\delta_x(x) = 1$ and $\delta_x(y) = 0$ if $y \neq x$.

Let μ be a measure on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. We denote its l -fold convolution by

$$\mu^{*(l)} = \underbrace{\mu * \dots * \mu}_l.$$

We denote the convolution of μ with a measure ν on \mathbb{F}_p^d by

$$[\mu.\nu](x) = \sum_{g \in G} \mu(g) \nu(g^{-1}.x)$$

which is a measure on \mathbb{F}_p^d .

The left regular representation on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ is denoted by \mathcal{L} and the representation obtained by composing the homomorphism θ with the left regular representation of $\mathrm{SL}_d(\mathbb{F}_p)$ is denoted by \mathcal{L}^θ . They are defined by the formulas

$$[\mathcal{L}(g)f](h) = f(g^{-1}h) \quad \text{and} \quad [\mathcal{L}^\theta(g)f'](\sigma) = f'(\theta(g)^{-1}\sigma)$$

for $f \in \mathbb{C}[\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)]$, $f' \in \mathbb{C}[\mathrm{SL}_d(\mathbb{F}_p)]$, $g, h \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ and $\sigma \in \mathrm{SL}_d(\mathbb{F}_p)$. In addition, we denote by \mathcal{L}_0 and \mathcal{L}_0^θ the restrictions of \mathcal{L} and \mathcal{L}^θ to the corresponding codimension one subspaces orthogonal to the constants.

Let π be a representation of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ and μ a probability measure on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. We write

$$\pi(\mu) = \sum_{g \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)} \mu(g) \pi(g)$$

which is an operator acting on the relevant representation space. Compare these with the definition of $\mathcal{L}(G, S)$ in the previous section.

We can generalize the notion of the random walk by considering random elements X_l having an arbitrary common law μ instead of a uniform distribution on a finite set S . Note that with the above notation, the law of Y_l is the probability measure

$$\mathcal{L}(\mu)^l \delta_1 = \mu^{*(l)}.$$

Now we state a more general version of Theorem 1 with a slight improvement in the bound.

Theorem 2. *There is a constant c depending only on d such that the following holds. Let μ be a probability measure on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$. Let α be the maximal probability for the event that a random element*

$X \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ of law μ takes a given point $x \in \mathbb{F}_p^d$ to a given point $y \in \mathbb{F}_p^d$. That is

$$\alpha = \max_{x, y \in \mathbb{F}_p^d} \mu.\delta_x(y).$$

Then

$$1 - \|\mathcal{L}_0(\mu)\| \geq c \min\{1 - \|\mathcal{L}_0^\theta(\mu)\|, 1 - \alpha\}.$$

In the setting of Theorem 1, for every $x, y \in \mathbb{F}_p$ there is at least one element $g \in S$ such that $g.x \neq y$. Indeed, in the opposite case, S would be contained in a coset of a subgroup isomorphic to $\mathrm{SL}_d(\mathbb{F}_p)$. Thus $\alpha \leq 1 - |S|^{-1}$, and Theorem 2 indeed contains Theorem 1 as a special case. In the rest of the paper we prove Theorem 2.

3. NON-CONCENTRATION ON SUBGROUPS

In this section, we make stronger assumptions on μ than in Theorem 2, but we will see in Section 4.2 that the general case can be reduced to this one. Let $v_0 \in \mathbb{F}_p^d$ be an arbitrary point, and consider the sequence of probability measures $\eta_l = \mu^{*(l)}.\delta_{v_0}$. These measures can be thought of as the laws of the steps of a random walk on \mathbb{F}_p^d starting from the point v_0 , and the steps being made by applying a random element of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ with law μ .

The purpose of this section is to prove the following result.

Proposition 3. *Suppose that μ is symmetric and*

$$\|\mu.\delta_x\|_{L^2} \leq \frac{3}{4}$$

for all $x \in \mathbb{F}_p^d$. Suppose further that

$$\|\mathcal{L}_0^\theta(\mu)\| \leq \frac{1}{2}.$$

Then for $l = \lfloor 2^{15} d \log p \rfloor$, we have

$$\|\eta_l\|_{L^\infty} \leq \|\eta_l\|_{L^2} \leq 4p^{-d/4}.$$

One can interpret this proposition as a non-concentration estimate. Indeed, the set $A_{v_0, u_0} \subset \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ consisting of elements g with the property $g.v_0 = u_0$ is a coset of a subgroup isomorphic to $\mathrm{SL}_d(\mathbb{F}_p)$. Moreover,

$$\mathbb{P}(Y_l \in A_{v_0, u_0}) = \mu^{*(l)}(A_{v_0, u_0}) = \eta_l(u_0),$$

which is estimated in the proposition.

We keep all constants explicit in this section for the sake of clarity, but we make no efforts to optimize them.

3.1. Some properties of the Fourier transform. We introduce some notation related to the Fourier transform on \mathbb{F}_p^d and some conventions for normalization. We denote the vector space of complex valued functions on \mathbb{F}_p^d by $\mathbb{C}[\mathbb{F}_p^d]$. Let $f \in \mathbb{C}[\mathbb{F}_p^d]$ and define its Fourier transform by

$$\widehat{f}(\xi) = \sum_{x \in \mathbb{F}_p^d} e(\langle x, \xi \rangle) f(x),$$

where $e(y) = e^{-2\pi i y/p}$ for $y \in \mathbb{Z}/p\mathbb{Z} = \mathbb{F}_p$.

We distinguish the space on which the Fourier transform is defined denoting it by $\widehat{\mathbb{F}}_p^d$. This space is of course isomorphic to \mathbb{F}_p^d , but we make this distinction in our notation because it will be convenient for us to use different normalizations for the L^q norms of functions on these two spaces. For functions $f \in \mathbb{C}[\mathbb{F}_p^d]$ and $\varphi \in \mathbb{C}[\widehat{\mathbb{F}}_p^d]$, we define these norms by

$$\|f\|_{L^q} := \left(\sum_{x \in \mathbb{F}_p^d} f(x)^q \right)^{1/q} \quad \text{and} \quad \|\varphi\|_{\widehat{L}^q} := \left(\frac{1}{p^d} \sum_{\xi \in \widehat{\mathbb{F}}_p^d} \varphi(\xi)^q \right)^{1/q}.$$

With this normalization, Plancherel's formula becomes $\|f\|_{L^2} = \|\widehat{f}\|_{\widehat{L}^2}$.

Remove the origin 0 from the set $\widehat{\mathbb{F}}_p^d$ and denote it by X . Let ι be the extension map from functions on X to functions on $\widehat{\mathbb{F}}_p^d$, i.e. $\iota(f)[x] = f(x)$ for $x \in \widehat{\mathbb{F}}_p^d \setminus 0$ and $\iota(f)[0] = 0$, and ι^* the natural projection from functions on $\widehat{\mathbb{F}}_p^d$ to functions on X (we will also use ι, ι^* to denote the exact same maps between $L^q(\widehat{\mathbb{F}}_p^d \setminus 0)$ and $L^q(\widehat{\mathbb{F}}_p^d)$). The norm \widehat{L}^q on X is defined as $\|\varphi\|_{\widehat{L}^q} = \|\iota(f)\|_{\widehat{L}^q}$, i.e. the total mass of the measure on X with respect to which the \widehat{L}^q -norms are defined is $(p^d - 1)/p^d$.

Let \mathcal{A} denote the natural actions of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ on both L^q and \widehat{L}^q , and \mathcal{A}^θ the corresponding action of the linear part of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ on \widehat{L}^q . Note that for any $f \in \widehat{L}^q$, $g \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$, and $x \in \mathbb{F}_p^d$

$$|(\mathcal{A}(g)f)(x)| = (\mathcal{A}^\theta(g)|f|)(x).$$

It would sometimes be useful to think of \mathcal{A}^θ also as an action of the group $\mathrm{SL}_d(\mathbb{F}_p)$.

The purpose of this section is to prove the following estimate, showing that if η is a probability measure on \mathbb{F}_p^d whose L^2 -norm is not too concentrated at a single atom and the nontrivial Fourier coefficients of η are not very small then the absolute value of the nontrivial Fourier coefficients of η cannot be almost constant.

Proposition 4. *Let η be a probability measure on \mathbb{F}_p^d . Suppose that*

$$\eta(x) \leq \frac{40}{41} \|\eta\|_{L^2} \quad \text{for every } x \in \mathbb{F}_p^d$$

and

$$(2) \quad \|\widehat{\eta}\|_{\widehat{L}^4} \geq 4p^{-d/4}.$$

Then there is an $h \in \text{SL}_d(\mathbb{F}_p)$ so that

$$\|(|\widehat{\eta}| - \mathcal{A}^\theta(h)|\widehat{\eta}|)\|_{\widehat{L}^4} \geq \frac{7}{100} \|\widehat{\eta}\|_{\widehat{L}^4}.$$

A key ingredient is the use of Plancherel's formula for the measure $\eta * \check{\eta}$, where $\check{\eta}$ denotes the probability measure on \mathbb{F}_p^d given by the formula:

$$\check{\eta}(x) = \eta(-x).$$

Observe that the Fourier transform of $\eta * \check{\eta}$ is $|\widehat{\eta}|^2$ and we need to show that it is not too close to constant. This is equivalent to $\eta * \check{\eta}$ being far from a Dirac measure supported at 0. This latter property of $\eta * \check{\eta}$ is proved in the next Lemma.

Lemma 5. *Let η be a probability measure on \mathbb{F}_p^d , and suppose that*

$$\eta(x) \leq \frac{40}{41} \|\eta\|_{L^2} \quad \text{for every } x \in \mathbb{F}_p^d.$$

Then

$$\|\iota^*(\eta * \check{\eta})\|_{L^2}^2 \geq \frac{1}{42} \|\eta * \check{\eta}\|_{L^2}^2.$$

Proof. By simple calculation:

$$\begin{aligned} \|\iota^*(\eta * \check{\eta})\|_{L^2}^2 &\equiv \sum_{x \neq 0} (\eta * \check{\eta})(x)^2 = \sum_{x \neq 0} \left[\sum_{y, z: y-z=x} \eta(y)\eta(z) \right]^2 \\ &\geq \sum_{x \neq 0} \sum_{y, z: y-z=x} [\eta(y)\eta(z)]^2 = \sum_{y, z: y \neq z} \eta(y)^2 \eta(z)^2 \\ (3) \quad &= \frac{1}{2} \left[\left(\sum_{y \in \mathbb{F}_p^d} \eta(y)^2 \right)^2 - \sum_{y \in \mathbb{F}_p^d} \eta(y)^4 \right] \end{aligned}$$

Using the assumption in the lemma:

$$(4) \quad \sum_{y \in \mathbb{F}_p^d} \eta(y)^4 \leq \max_{y \in \mathbb{F}_p^d} \{\eta(y)^2\} \sum_{y \in \mathbb{F}_p^d} \eta(y)^2 \leq \frac{1600}{1681} \left[\sum_{y \in \mathbb{F}_p^d} \eta(y)^2 \right]^2$$

From inequalities (3) and (4) we get

$$\sum_{x \neq 0} (\eta * \check{\eta})(x)^2 \geq \frac{1}{42} \|\eta\|_{L^2}^4$$

because $1/42 < (1 - 1600/1681)/2$. Thus

$$\eta * \check{\eta}(0)^2 = \left[\sum_{x \in \mathbb{F}_p^d} \eta(x)^2 \right]^2 = \|\eta\|_{L^2}^4 \leq 42 \sum_{x \neq 0} (\eta * \check{\eta})(x)^2,$$

which implies the claim. \square

A tool which will help us relate the L^2 and L^4 norm is the Mazur map. We recall its definition and properties from the book of Benyamini and Lindenstrauss [BL00, Chapter 9]. We denote by $S(L^q)$ the unit sphere of the space $L^q(X)$. For $f \in S(L^4)$, the Mazur map is defined by

$$\phi(f) = |f|^2 \operatorname{sign}(f),$$

where $\operatorname{sign}(f) = f/|f|$ for $f \neq 0$ and 0 otherwise.

The Mazur map is a homeomorphism from $S(L^4)$ to $S(L^2)$. Moreover, we have the following inequalities.

Theorem C ([BL00, Theorem 9.1]). *For $f_1, f_2 \in S(L^4)$, we have*

$$\|f_1 - f_2\|_{L^4} \geq \frac{1}{2} \|\phi(f_1) - \phi(f_2)\|_{L^2}.$$

For proof, see [BL00, Proof of Theorem 9.1], applied to $p = 4$ and $q = 2$.

Proof of Proposition 4. The existence of a $h \in \operatorname{SL}_d(\mathbb{F}_p)$ as in the proposition follows from the following estimate:

$$\begin{aligned} (5) \quad & \frac{1}{\#\operatorname{SL}_d(\mathbb{F}_p)} \sum_{h \in \operatorname{SL}_d(\mathbb{F}_p)} \frac{\|(|\hat{\eta}| - \mathcal{A}^\theta(h) |\hat{\eta}|)\|_{\widehat{L}^4}}{\|\hat{\eta}\|_{\widehat{L}^4}} \\ & \geq \frac{1}{2\#\operatorname{SL}_d(\mathbb{F}_p)} \sum_{h \in \operatorname{SL}_d(\mathbb{F}_p)} \frac{\|(|\hat{\eta}|^2 - \mathcal{A}^\theta(h) |\hat{\eta}|^2)\|_{\widehat{L}^2}}{\| |\hat{\eta}|^2 \|_{\widehat{L}^2}} \\ & \geq \frac{\| |\hat{\eta}|^2 - \frac{1}{\#\operatorname{SL}_d(\mathbb{F}_p)} \sum_{h \in \operatorname{SL}_d(\mathbb{F}_p)} \mathcal{A}^\theta(h) |\hat{\eta}|^2 \|_{\widehat{L}^2}}{2 \| |\hat{\eta}|^2 \|_{\widehat{L}^2}} \\ & \geq \frac{\|\iota^*(|\hat{\eta}|^2) - c\|_{\widehat{L}^2(X)}}{2 \| |\hat{\eta}|^2 \|_{\widehat{L}^2}} \end{aligned}$$

for $c = (p^d - 1)^{-1} \sum_{x \neq 0} |\widehat{\eta}(x)|^2$; in particular, $0 \leq c \leq 1$. Note that Theorem C was used to pass from the first to the second line in (5).

Since $\widehat{\eta}(0) = 1$, and using (2), we have that

$$\begin{aligned} \|\iota^*(|\widehat{\eta}|^2) - c\|_{\widehat{L}^2(X)}^2 &= \| |\widehat{\eta}|^2 - c \|_{\widehat{L}^2}^2 - p^{-d}(|\widehat{\eta}(0)|^2 - c)^2 \\ &\geq \| |\widehat{\eta}|^2 - c \|_{\widehat{L}^2}^2 - p^{-d} \\ &= \|\eta * \check{\eta} - c\delta_0\|_{L^2}^2 - p^{-d} \\ &\geq \|\iota^*(\eta * \check{\eta})\|_{L^2}^2 - p^{-d} \\ &\geq \left(\frac{1}{42} - \frac{1}{256}\right) \|\eta * \check{\eta}\|_{L^2}^2 \end{aligned}$$

since $\|\eta * \check{\eta}\|_{L^2} = \| |\widehat{\eta}|^2 \|_{\widehat{L}^2} = \|\widehat{\eta}\|_{\widehat{L}^4}^2 \geq 16p^{-d/2}$. In the last line, we used Lemma 5. Using equation (5) we can now conclude that there is some h so that

$$\frac{\|(|\widehat{\eta}| - \mathcal{A}^\theta(h) |\widehat{\eta}|)\|_{\widehat{L}^4}}{\|\widehat{\eta}\|_{\widehat{L}^4}} \geq \frac{\sqrt{1/42 - 1/256}}{2} \geq \frac{7}{100}$$

establishing the proposition. \square

3.2. A consequence of the Riesz-Thorin interpolation theorem. We will now use the Riesz Thorin interpolation theorem to study how $\mathcal{A}^\theta(\mu)$ acts on the Fourier transform of measures on \mathbb{F}_p^d with respect to the \widehat{L}^4 -norm.

Proposition 6. *Let μ be a probability measure on $\mathrm{SL}_d(\mathbb{F}_p)$ satisfying the conditions of Proposition 3 and η a probability measure on \mathbb{F}_p^d satisfying the conditions in Proposition 4. Then*

$$\|\mathcal{A}^\theta(\mu^{*(5)}) |\widehat{\eta}|\|_{\widehat{L}^4} \leq e^{-5 \cdot 2^{-14}} \|\widehat{\eta}\|_{\widehat{L}^4}.$$

Lemma 7. *Let f, g be nonnegative functions on a σ -finite measure space. Then*

$$\frac{1}{2}(\|f\|_{L^4}^4 + \|g\|_{L^4}^4) \geq \left\| \frac{f+g}{2} \right\|_{L^4}^4 + 7 \left\| \frac{f-g}{2} \right\|_{L^4}^4.$$

Proof. This follows easily from the inequality

$$\frac{1+x^4}{2} \geq \left(\frac{1+x}{2}\right)^4 + 7 \left(\frac{1-x}{2}\right)^4$$

which is valid for $x \geq 0$. \square

Proof of Proposition 6. Consider for $h \in \mathrm{SL}_d(\mathbb{F}_p)$ the operator

$$(\mathcal{A}^\theta(h) - 1)\mathcal{A}^\theta(\mu^{*(10)}),$$

where 1 denotes the identity operator. The assumption $\|\mathcal{L}_0^\theta(\mu)\|_{L^2} \leq \frac{1}{2}$ implies $\|\mathcal{A}_0^\theta(\mu)\|_{L^2} \leq \frac{1}{2}$. Since $\mathcal{A}^\theta(h) - 1$ annihilate the constants and it has L^2 and L^∞ norm at most 2, we have that

$$\begin{aligned} \|(\mathcal{A}^\theta(h) - 1)\mathcal{A}^\theta(\mu^{*(10)})\|_{\widehat{L}^2} &\leq 2^{-9} \\ \|(\mathcal{A}^\theta(h) - 1)\mathcal{A}^\theta(\mu^{*(10)})\|_{\widehat{L}^\infty} &\leq 2. \end{aligned}$$

Hence by interpolation

$$\|(\mathcal{A}^\theta(h) - 1)\mathcal{A}^\theta(\mu^{*(10)})\|_{\widehat{L}^4} \leq 2^{-4}.$$

For any $h \in \mathrm{SL}_d(\mathbb{F}_p)$,

$$\begin{aligned} \|\mathcal{A}^\theta(h) |\widehat{\eta}| - |\widehat{\eta}|\|_{\widehat{L}^4} &\leq \left\| \mathcal{A}^\theta(h)\mathcal{A}^\theta(\mu^{*(10)}) |\widehat{\eta}| - \mathcal{A}^\theta(\mu^{*(10)}) |\widehat{\eta}| \right\|_{\widehat{L}^4} + \\ &\quad + 2 \|\mathcal{A}^\theta(\mu^{*(10)}) |\widehat{\eta}| - |\widehat{\eta}|\|_{\widehat{L}^4} \\ &\leq 2^{-4} \|\widehat{\eta}\|_{\widehat{L}^4} + 2 \|\mathcal{A}^\theta(\mu^{*(10)}) |\widehat{\eta}| - |\widehat{\eta}|\|_{\widehat{L}^4} \end{aligned}$$

hence using Proposition 4 there is a $h \in \mathrm{SL}_d(\mathbb{F}_p)$ so that

$$(6) \quad \|\mathcal{A}^\theta(\mu^{*(10)}) |\widehat{\eta}| - |\widehat{\eta}|\|_{\widehat{L}^4} \geq \frac{1}{2} \left(\frac{7}{100} - \frac{1}{16} \right) \|\widehat{\eta}\|_{\widehat{L}^4} = \frac{3}{800} \|\widehat{\eta}\|_{\widehat{L}^4}.$$

As μ is symmetric,

$$\begin{aligned} (7) \quad &\|\mathcal{A}^\theta(\mu^{*(10)}) |\widehat{\eta}| - |\widehat{\eta}|\|_{\widehat{L}^4} \\ &\leq \sum_{g, g'} \mu^{*(5)}(g) \mu^{*(5)}(g') \|(\mathcal{A}^\theta(g) - \mathcal{A}^\theta(g')) |\widehat{\eta}|\|_{\widehat{L}^4}. \end{aligned}$$

But by Lemma 7, for any nonnegative $f \in \widehat{L}^4$

$$\left\| \frac{(\mathcal{A}^\theta(g) + \mathcal{A}^\theta(g'))f}{2} \right\|_{\widehat{L}^4}^4 \leq \|f\|_{\widehat{L}^4}^4 - \frac{7}{16} \|(\mathcal{A}^\theta(g) - \mathcal{A}^\theta(g'))f\|_{\widehat{L}^4}^4,$$

hence as $(1 - x)^{1/4} \leq 1 - x/4$ for $0 \leq x \leq 1$

$$(8) \quad \left\| \frac{(\mathcal{A}^\theta(g) + \mathcal{A}^\theta(g'))f}{2} \right\|_{\widehat{L}^4} \leq \|f\|_{\widehat{L}^4} - \frac{7}{64} \|(\mathcal{A}^\theta(g) - \mathcal{A}^\theta(g'))f\|_{\widehat{L}^4}.$$

Applying (6), (7) and (8) it follows that

$$\begin{aligned}
\|\mathcal{A}^\theta(\mu^{*(5)})|\widehat{\eta}|\|_{\widehat{L}^4} &\leq \sum_{g,g'} \mu^{*(5)}(g)\mu^{*(5)}(g') \left\| \left(\frac{\mathcal{A}^\theta(g) + \mathcal{A}^\theta(g')}{2} \right) |\widehat{\eta}| \right\|_{\widehat{L}^4} \\
&\leq \|\widehat{\eta}\|_{\widehat{L}^4} - \frac{7}{64} \sum_{g,g'} \mu^{*(5)}(g)\mu^{*(5)}(g') \|\mathcal{A}^\theta(g) - \mathcal{A}^\theta(g')\| |\widehat{\eta}|_{\widehat{L}^4} \\
&\leq \left(1 - \frac{21}{64 \cdot 800} \right) \|\widehat{\eta}\|_{\widehat{L}^4} \\
&\leq e^{-5 \cdot 2^{-14}} \|\widehat{\eta}\|_{\widehat{L}^4}.
\end{aligned}$$

□

3.3. Finishing the proof of Proposition 3. We consider two cases. First we suppose that η_l does not concentrate too big mass on a single atom, that is:

$$(9) \quad \eta_l(x) \leq \frac{40}{41} \|\eta_l\|_{L^2}$$

for all $x \in \mathbb{F}_p^d$. If this is the case, we can apply Proposition 6 and conclude either

$$\|\widehat{\eta}_{l+5}\|_{\widehat{L}^4} \leq \|\mathcal{A}^\theta(\mu^{*(5)})|\widehat{\eta}|\|_{\widehat{L}^4} \leq e^{-5 \cdot 2^{-14}} \|\widehat{\eta}\|_{\widehat{L}^4}.$$

or $\|\widehat{\eta}_l\|_{\widehat{L}^4} \leq 4p^{-d/4}$.

On the other hand, we always have the trivial inequality $\|\widehat{\eta}_{l+1}\|_{\widehat{L}^4} \leq \|\widehat{\eta}_l\|_{\widehat{L}^4}$. Thus if we have $\|\widehat{\eta}_k\|_{\widehat{L}^4} \geq 4p^{-d/4}$ for some integer $k > 0$, then there are at most

$$2^{14}d \log p$$

many nonnegative integers $l < k$ such that (9) holds.

Now we turn to the second case, when (9) does not hold, that is, there is a point $x_0 \in \mathbb{F}_p^d$ such that

$$\eta_l(x_0) \geq \frac{40}{41} \|\eta_l\|_{L^2}.$$

In this case η_l is very close to a constant multiple of δ_{x_0} in the L^2 norm so we can estimate $\|\eta_{l+1}\|_{L^2}$ using the assumption $\|\mu \cdot \delta_{x_0}\|_{L^2} \leq 3/4$.

More precisely, we can write

$$\begin{aligned}
\|\eta_{l+1}\|_{L^2} &= \|\mu \cdot \eta_l\|_{L^2} \\
&\leq \frac{3}{4} \eta_l(x_0) + \sqrt{\|\eta_l\|_{L^2}^2 - \eta_l^2(x_0)} \\
&\leq \left(\frac{3}{4} + \frac{9}{41} \right) \|\eta_l\|_{L^2} < e^{-2^{-6}} \|\eta_l\|_{L^2}.
\end{aligned}$$

Since η_k is a probability measure for all integers $k \geq 0$, we have $\|\eta_k\|_{L^2} \geq p^{-d/2}$. Therefore it follows that the number of nonnegative integers $l < k$ such that (9) fails is at most

$$2^5 d \log p.$$

If we combine this with the estimate for the number of steps when (9) holds, we can conclude that for

$$k = \lfloor 2^{15} d \log p \rfloor$$

we have $\|\widehat{\eta}_k\|_{\widehat{L}^4} \leq 4p^{-d/4}$, hence $\|\eta_k\|_{L^2} = \|\widehat{\eta}_k\|_{\widehat{L}^2} \leq 4p^{-d/4}$.

4. THE BOURGAIN-GAMBURD METHOD

We use the method of Bourgain and Gamburd to prove Theorem 2. This material is fairly standard now, and most ideas have already appeared in earlier works. This method proves that a random walk has spectral gap if two conditions hold. First, the group G should have no low dimensional representations. Second, if there is a subset A of G of size approximately $|G|^\beta$ that does not grow under multiplication, then the probability that the random walk hits this set after approximately $l = \log |G|$ steps should be very small, e.g. $|G|^\varepsilon$.

The method uses the notion of product sets, which we define now. Let $A \subset G$ be a set; its l -fold product set is the set

$$\Pi_l A = \{a_1 \cdots a_l : a_1, \dots, a_l \in A\}.$$

The method can be summarized in the following theorem.

Theorem D (Bourgain, Gamburd). *There is an absolute constant C , and for any $\varepsilon > 0$ there is a $\delta > 0$ such that the following holds. Let G be a finite group and π an irreducible unitary representation of it. Let μ be a symmetric probability measure on G . Let $l_1 > 0$ be an integer and suppose that for any symmetric set $A \subset G$ that satisfies*

$$(10) \quad \mu^{*(l)}(A) \geq |G|^{-\varepsilon}$$

for some integer $l \geq l_1$, we either have

$$(11) \quad |\Pi_3 A| \geq |G|^\varepsilon \cdot |A| \quad \text{or} \quad |A| \geq (\dim \pi)^{-1/3} |G|.$$

Then

$$\|\pi(\mu)\| < (C \dim \pi)^{-\delta/l_1}.$$

Note that if ε is too large or $\dim \pi$ too small there may be no probability measure μ satisfying the conditions of the theorem. Indeed, if $\lfloor (\dim \pi)^{-1/3} |G| \rfloor > |G|^{1-\varepsilon}$ then for any probability measure μ we can find a set A with

$$(\dim \pi)^{-1/3} |G| > |A| > |G|^{1-\varepsilon}$$

so that $\mu(A) > |G|^{-\varepsilon}$, violating the conditions of the theorem since clearly $|\Pi_3 A| \leq |G|$.

This theorem is implicitly contained in the paper [BG08], and variants of its proof appeared in many papers. In particular, [Kow13, Corollary 4.4] contains a version with explicit constants, but unfortunately, as it is stated that version only applies to groups without large normal subgroups. For completeness, we include the proof of Theorem D in Section 4.3.

We comment on the role of the triple product set $\Pi_3 A$ in the theorem. The following Lemma shows that if the set $\Pi_k A$ is much larger than A , then so is $\Pi_3 A$. Hence an equivalent theorem could be stated with $\Pi_k A$ instead of $\Pi_3 A$ for any integer $k \geq 3$. This observation will be important for us, since in our proof of (11), we will estimate the size of $\Pi_{29} A$. Note, however, that for nonabelian groups it may well happen that $\Pi_2 A$ is of comparable size to A but $\Pi_3 A$ is much bigger. The lemma below (in a less explicit form) is due to Tao [Tao08, Lemma 3.4]; in this form, which is a simple corollary of the Ruzsa Triangle Inequality (cf. [Tao08, Lemma 3.2] and the references given there), it can be found e.g. in [LV].

Lemma E. *Let $A \subset G$ be a symmetric subset of a group. Then for any integer $k \geq 3$, we have*

$$\frac{|\Pi_k A|}{|A|} \leq \left(\frac{|\Pi_3 A|}{|A|} \right)^{k-2}.$$

The following lemma gives a lower bound on the dimension of non-trivial representations of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$.

Lemma F (Landazuri, Seitz). *If π is a nontrivial representation of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$, then*

$$\dim \pi \geq \begin{cases} \frac{1}{2}(p-1) & \text{if } d = 2 \\ p^{d-1} - 1 & \text{otherwise.} \end{cases}$$

Proof. Since $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ is generated by subgroups isomorphic to $\mathrm{SL}_d(\mathbb{F}_p)$, the restriction of π to one of these must be non-trivial. Then the bound claimed in the lemma is in [LS74, p. 419]. \square

In Section 4.1, we show that if μ is a measure that satisfies the conditions in Proposition 3, then any set A that satisfies (10), also satisfies the growth condition (11).

In Section 4.2, we construct a measure μ_0 which satisfies the conditions in Proposition 3 using the measure μ from Theorem 2. We

will relate the random walks generated by the measures μ and μ_0 and conclude the proof of Theorem 2.

4.1. Growth of product sets. The following result is not new, a version with different constants could be deduced from the more general results [PS10, Theorem 7] or [SGV12, Proposition 27]. Since this special case is much simpler, we provide a quick proof for completeness.

Proposition 8. *Let μ be a symmetric probability measure on G that satisfies the conditions required in Proposition 3. Then there is an $\varepsilon > 0$ depending only on d , such that the following holds. Let $A \subset \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ be a symmetric set that satisfies*

$$(12) \quad \mu^{*(l)}(A) \geq |\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)|^{-\varepsilon}$$

for some integer $l \geq l_1 = \lfloor 2^{15} d^2 \log p \rfloor$. Then

$$\Pi_{29} A = \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p).$$

In what follows, we assume that μ satisfies the conditions of Proposition 3 and $A \subset \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ is a set that satisfies the conditions of Proposition 8.

We first show that $\Pi_3 A$ projects onto $\mathrm{SL}_d(\mathbb{F}_p)$. To this end, we exploit the assumption of the spectral gap in the quotient. Then we show that there is a pure translation in $\Pi_7 A$. We conjugate this with elements of $\Pi_3 A$, to get all pure translations in $\Pi_{26} A$. Finally we multiply this with $\Pi_3 A$ to recover the whole group.

The same strategy was employed in [SGV12], but our proof differs in the way we produce the first pure translation (proof of Lemma 10 below). In [SGV12] the inequality $|\Pi_4 A| > |\Pi_3 A|$ was exploited (this inequality holds if one knows as is the case in [SGV12] that A is generating, unless of course if $\Pi_3 A$ is already everything), which implies that $\Pi_4 A$ must contain two elements with the same linear part. In the present paper, we give a different proof based on an averaging argument, which works well in the continuous setting of [LV14b] as well.

Lemma 9. *We have $\theta(\Pi_3 A) = \mathrm{SL}_d(\mathbb{F}_p)$.*

Proof. We will show that

$$|\theta(A)| \geq \frac{|\mathrm{SL}_d(\mathbb{F}_p)|}{D^{1/3}},$$

where D is the minimal dimension of a non-trivial representation of $\mathrm{SL}_d(\mathbb{F}_p)$. Then the claim $\Pi_3 \theta(A) = \mathrm{SL}_d(\mathbb{F}_p)$ follows from a theorem of Nikolov and Pyber [NP11, Corollary 1] (based on a paper of Gowers [Gow08]).

We begin by noting the identity

$$\theta(\mu^{*(l)}) = \mathcal{L}^\theta(\mu)^l \delta_1.$$

By the assumption in Proposition 3, we have $\|\mathcal{L}_0^\theta(\mu)\| \leq 1/2$. We can write $\delta_1 = \varphi_1 + \varphi_2$, such that $\varphi_1 \equiv 1/|\mathrm{SL}_d(\mathbb{F}_p)|$, and φ_2 is orthogonal to the constant. Then

$$\|\theta(\mu^{*(l)})\|_2 \leq |\mathrm{SL}_d(\mathbb{F}_p)|^{-1/2} + \frac{1}{2^l} \leq 2|\mathrm{SL}_d(\mathbb{F}_p)|^{-1/2},$$

since $l \geq l_1 \geq d^2 \log p / \log 2$.

By the assumption in Proposition 8, we have

$$\sum_{g \in A} \mu^{*(l)}(g) \geq |\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)|^{-\varepsilon}.$$

By the Cauchy-Schwartz inequality,

$$\begin{aligned} |\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)|^{-\varepsilon} &\leq \sum_{\sigma \in \theta(A)} \theta(\mu^{*(l)})(\sigma) \\ &\leq |\theta(A)|^{1/2} \|\theta(\mu^{*(l)})\|_2. \end{aligned}$$

Combining with the inequality in the previous paragraph, this implies

$$|\theta(A)| \geq \frac{|\mathrm{SL}_d(\mathbb{F}_p)|}{4|\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)|^{2\varepsilon}}.$$

To finish, we note that any non-trivial representation of $\mathrm{SL}_d(\mathbb{F}_p)$ is of dimension $\geq (p^{d-1} - 1)/2$ (see Lemma F), and $|\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)| \leq p^{d^2+d}$. Now the lemma follows from the remarks at the beginning of the proof, if ε is sufficiently small depending on d . If p is sufficiently large, any $\varepsilon \leq 1/(6d+6)$ works. \square

Lemma 10. *There is a non-zero pure translation in $\Pi_7 A$, that is, there is an element $g_0 \in \Pi_7 A$, such that $\theta(g_0) = 1$ and $v(g_0) \neq 0$.*

Proof. By Lemma 9, there is a map $F : \mathrm{SL}_d(\mathbb{F}_p) \rightarrow \Pi_3 A$ such that $\sigma = \theta(F(\sigma))$ for all $\sigma \in \mathrm{SL}_d(\mathbb{F}_p)$. We define

$$v_0 = \sum_{\sigma \in \mathrm{SL}_d(\mathbb{F}_p)} v(F(\sigma)).$$

We show that v_0 is not a fixed point for all elements of A under the natural action. To this end, we write

$$\eta_l = \mu^{*(l)} \cdot \delta_{v_0}.$$

Since $l \geq 2^{15} d \log p$, we can apply Proposition 3, and we have

$$\|\eta_l\|_{L^\infty} \leq \|\eta_l\|_{L^2} \leq 4p^{-d/4}.$$

Denoting by $G_{v_0} \subset \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ the stabilizer of the point $v_0 \in \mathbb{F}_p^d$, we have

$$\mu^{*(l)}(G_{v_0}) = \eta_l(v_0) \leq 4p^{-d/4}.$$

If ε is small enough, the assumption in Proposition 8 implies that $A \not\subset G_{v_0}$. That is, there is $g_1 \in A$ such that $g_1.v_0 \neq v_0$ as claimed.

We look at elements of the following form:

$$F_2(\sigma) = F(\theta(g_1)\sigma)^{-1}g_1F(\sigma) \in \Pi_7A.$$

By the definition of F , we have

$$\theta(F_2(\sigma)) = (\theta(g_1)\sigma)^{-1}\theta(g_1)\sigma = 1$$

for all $\sigma \in \mathrm{SL}_d(\mathbb{F}_p)$.

On the other hand

$$v(F_2(\sigma)) = F_2(\sigma).0 = -\sigma^{-1}\theta(g_1)^{-1}.v(F(\theta(g_1)\sigma)) + \sigma^{-1}\theta(g_1)^{-1}g_1.v(F(\sigma)).$$

(To see this, recall formula (1) for the inverse of an element of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$.) Then

$$\theta(g_1)\sigma.v(F_2(\sigma)) = -v(F(\theta(g_1)\sigma)) + g_1.v(F(\sigma)).$$

Since left multiplication by $\theta(g_1)$ is a permutation on $\mathrm{SL}_d(\mathbb{F}_p)$, we get

$$\begin{aligned} (13) \quad & \sum_{\sigma \in \mathrm{SL}_d(\mathbb{F}_p)} \theta(g_1)\sigma.v(F_2(\sigma)) \\ &= \sum_{\sigma \in \mathrm{SL}_d(\mathbb{F}_p)} [-v(F(\theta(g_1)\sigma)) + g_1.v(F(\sigma))] \\ (14) \quad &= -v_0 + g_1.v_0. \end{aligned}$$

If $v(F_2(\sigma))$ were 0 for all σ , then (13) would be 0. On the other hand (14) is clearly non-zero, by the choice of g_1 . This proves that for some choice of $\sigma \in \mathrm{SL}_d(\mathbb{F}_p)$, $g_0 = F_2(\sigma)$ satisfies the claims of the lemma. \square

Proof of Proposition 8. We consider the element $g_0 \in \Pi_7A$ found in Lemma 10 and all elements of the form $gg_0g^{-1} \in \Pi_{13}A$ for $g \in \Pi_3A$. Since $\theta(gg_0g^{-1}) = \theta(g)\theta(g)^{-1} = 1$, all of these are pure translations. On the other hand, $v(gg_0g^{-1}) = \theta(g).v(g_0)$, and $\theta(\Pi_3A) = \mathrm{SL}_d(\mathbb{F}_p)$, hence $\Pi_{13}A$ contains all non-zero pure translations. Then it follows that $\Pi_{26}A$ contains all pure translations.

Therefore, for a fixed $g \in \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$, the set $\Pi_{26}A \cdot g$ contains all elements of $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ whose linear part is $\theta(g)$. Since $\theta(\Pi_3A) = \mathrm{SL}_d(\mathbb{F}_p)$, $\Pi_{29}A = \mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$, as claimed. \square

4.2. Proof of Theorem 2. We fix an integer

$$l_0 \geq \max \left\{ \frac{3}{1-\alpha}, \frac{\log 2}{2-2\|\mathcal{L}_0^\theta(\mu)\|} \right\}$$

and set

$$\mu_0 = (\check{\mu} * \mu)^{*(l_0)},$$

where $\check{\mu}$ is the measure on $\mathbb{F}_p^d \rtimes \mathrm{SL}_d(\mathbb{F}_p)$ defined by

$$\check{\mu}(g) = \mu(g^{-1}).$$

The next lemma shows that the conditions of Propositions 3 and 8 hold for μ_0 .

Lemma 11. *With the notations above, the following holds:*

$$\begin{aligned} \|\mathcal{L}_0^\theta(\mu_0)\| &\leq \frac{1}{2}, \quad \text{and} \\ \|\mu_0 \cdot \delta_x\|_{L^2} &\leq \frac{3}{4} \quad \text{for all } x \in \mathbb{F}_p^d. \end{aligned}$$

Proof. The first claim follows from

$$\|\mathcal{L}_0^\theta(\mu_0)\| = \|\mathcal{L}_0^\theta(\check{\mu} * \mu)\|^{l_0} = \|\mathcal{L}_0^\theta(\mu)\|^{2l_0} \leq e^{(\|\mathcal{L}_0^\theta(\mu)\|-1)2l_0}$$

and the assumption $l_0 \geq \log 2 / (2 - 2\|\mathcal{L}_0^\theta(\mu)\|)$.

We turn to the proof of the second claim. For $x, y \in \mathbb{F}_p^d$ and a positive integer l , we write

$$\alpha_l(x, y) = (\check{\mu} * \mu)^{*(l)} \cdot \delta_y(x).$$

This is the probability that the random walk on \mathbb{F}_p^d generated by $\check{\mu} * \mu$ started from y is at the point x after l steps. It is easy to verify the identity

$$\alpha_{l+1}(x, y) = \sum_{z \in \mathbb{F}_p^d} \alpha_1(x, z) \alpha_l(z, y).$$

Write

$$\alpha_l = \max_{x, y \in \mathbb{F}_p^d} \alpha_l(x, y),$$

and observe that $\alpha_1 \leq \alpha$. We claim that $\alpha_{l_0} \leq 9/16$.

To show this, write

$$\begin{aligned} \alpha_{l+1}(x, y) &= \sum_{z \in \mathbb{F}_p^d} \alpha_1(x, z) \alpha_l(z, y) \\ &\leq \max_z \alpha_1(x, z) \cdot \max_z \alpha_l(z, y) \\ &\quad + (1 - \max_z \alpha_1(x, z)) \cdot (1 - \max_z \alpha_l(z, y)). \end{aligned}$$

In the domain $1/2 \leq s \leq 1$, $1/2 \leq t \leq 1$, the function $st - (1-s)(1-t)$ is monotone increasing in both variables. Thus

$$(15) \quad \alpha_{l+1}(x, y) \leq \alpha_1 \alpha_l + (1 - \alpha_1)(1 - \alpha_l)$$

provided

$$(16) \quad \max_z \alpha_1(x, z) \geq \frac{1}{2} \quad \text{and} \quad \max_z \alpha_l(z, y) \geq \frac{1}{2}.$$

If (16) fails, then

$$\alpha_{l+1}(x, y) \leq \min\{\max_z \alpha_1(x, z), \max_z \alpha_l(z, y)\} \leq \frac{1}{2},$$

so in either case we get

$$\alpha_{l+1} \leq \max\{\alpha_1 \alpha_l + (1 - \alpha_1)(1 - \alpha_l), 1/2\}.$$

If $\alpha_l \leq 1/2$ for some $l \leq l_0$, then there is nothing to prove, so we assume this is not the case. We can then write

$$\left(\alpha_{l+1} - \frac{1}{2}\right) \leq \alpha_1 \left(\alpha_l - \frac{1}{2}\right) - \frac{1}{2}(1 - \alpha_1) + (1 - \alpha_1)(1 - \alpha_l) \leq \alpha_1 \left(\alpha_l - \frac{1}{2}\right)$$

for any $l < l_0$. By iteration, and using $\alpha_1 \leq \alpha$, we get

$$\alpha_{l_0} \leq \frac{1}{2} + e^{(\alpha-1)l_0}.$$

Since we took $l_0 \geq 3/(1 - \alpha)$, this implies $\alpha_{l_0} \leq 9/16$, as claimed.

To finish the proof of the second claim of the lemma, we observe that

$$\mu_0 \cdot \delta_x(y) = \alpha_{l_0}(y, x) \leq \alpha_{l_0} \leq \frac{9}{16}.$$

This implies

$$\|\mu_0 \cdot \delta_x\|_{L^2}^2 \leq \|\mu_0 \cdot \delta_x\|_{L^\infty} \cdot \|\mu_0 \cdot \delta_x\|_{L^1} \leq \frac{9}{16},$$

which was to be proved. \square

We are now in a position to finish the proof of Theorem 2. By Lemma 11, the conditions of Propositions 3 and 8 are satisfied for μ_0 . By these propositions and Lemma E, we can apply Theorem D with $l_1 = \lfloor 2^{15} d^2 \log p \rfloor$ and some $\varepsilon > 0$ small enough depending on d . Thus we can conclude for all irreducible representations of $\mathbb{F}_p^d \rtimes \text{SL}_d(\mathbb{F}_p)$ that

$$\|\pi(\mu_0)\| < (C \dim \pi)^{-\delta/l_1}$$

(δ depending on ε , hence on d). Then by Lemma F we have

$$\|\pi(\mu_0)\| < C_1^{1/\log p} e^{-2^{-15} \delta/d}$$

if π is non-trivial with C_1 depending only on d . If p is sufficiently large depending on the constants in the above inequality (hence only on d), then we can write

$$\|\pi(\mu_0)\| \leq e^{-c_d},$$

for some number $c_d > 0$ depending only on d . Since there are only finitely many not large enough primes, and the set of probability measures satisfying the conclusions of Lemma 11 is compact, the above inequality holds for all p for some number c_d .

Note that $\|\mathcal{L}_0(\mu_0)\|$ is the maximum of $\|\pi(\mu_0)\|$ for π running through the non-trivial irreducible representations. Thus

$$\|\mathcal{L}_0(\mu)\| = \|\mathcal{L}_0(\mu_0)\|^{\frac{1}{2l_0}} \leq e^{-\frac{c_d}{2l_0}},$$

which was to be proved.

4.3. Proof of Theorem D. We suppose that the assumptions of the theorem hold for some $G, \pi, \mu, \varepsilon, l_1$ and prove the conclusion for some C, δ .

The proof due to Bourgain and Gamburd consists of two parts. First, we consider the L^2 -norms $\|\mu^{*(l)}\|_2$ for $l \geq l_1$ and give improved bounds as l increases. Second, we exploit the fact that the eigenvalues of convolution operators on $L^2(G)$ have high multiplicities and hence we can get an estimate on them when $\|\mu^{*(l)}\|_2$ is close to the optimal bound, that is $|G|^{-1/2}$. This second idea goes back to Sarnak and Xue [SX91].

We recall the “ L^2 -flattening Lemma” of Bourgain and Gamburd. This appeared implicitly in [BG08], and it is an application of the Balog-Szemerédi-Gowers theorem combined with some results of Tao [Tao08]. We use the version in [Var12, Lemma 15].

Lemma G (Bourgain, Gamburd). *Let ν_1 and ν_2 be two probability measures on a finite group G and let $K > 2$ be a number. If*

$$\|\nu_1 * \nu_2\|_2 \geq \frac{\|\nu_1\|_2^{1/2} \|\nu_2\|_2^{1/2}}{K}$$

then there is a symmetric set $S \subset G$ with

$$(17) \quad \frac{1}{CK^C \|\nu_1\|_2^2} \leq |S| \leq \frac{CK^C}{\|\nu_1\|_2^2},$$

$$|\Pi_3 S| \leq CK^C |S|, \quad \min_{g \in S} (\check{\nu}_1 * \nu_1)(g) \geq \frac{1}{CK^C |S|},$$

where C is an absolute constant.

We now prove Theorem D. Fix a number K in such a way that $CK^C \leq |G|^\varepsilon$, where C is from Lemma G and ε is from Theorem D.

(We may assume that $|G|$ is larger than any absolute constant, since the theorem is vacuous when $|G|$ is small, if we set the constant C large enough in the theorem.)

By Lemma G, for all $l \geq l_1$, we have either $\|\mu^{*(2l)}\|_2 \leq \|\mu^{*(l)}\|_2/K$, or there is a symmetric set $S \subset G$ such that $|\Pi_3 S| \leq |G|^\varepsilon |S|$ and

$$\mu^{*(2l)}(S) = \check{\mu}^{*(l)} * \mu^{*(l)}(S) \geq \frac{|S|}{CK^C|S|} \geq |G|^{-\varepsilon}.$$

In the latter case, S satisfies condition (10) in the theorem and fails the first alternative of (11). Then we must have $|S| \geq (\dim \pi)^{-1/3} |G|$ in this case and hence by (17)

$$\|\mu^{*(l)}\|_2^2 \leq |G|^\varepsilon (\dim \pi)^{1/3} |G|^{-1}.$$

We have already noted that the assumption of Theorem D may hold only if $|G|^\varepsilon \leq 2(\dim \pi)^{1/3}$. Hence in this second case we must have that $\|\mu^{*(l)}\|_2^2 \leq 2(\dim \pi)^{2/3} |G|^{-1}$.

We consider the sequence $a_k := \|\mu^{*(2^k l_1)}\|_2^2$ for $k = 0, 1, \dots$. The argument of the previous paragraph shows that either $a_{k+1} \leq a_k/K^2$ or $a_k \leq 2(\dim \pi)^{2/3} |G|^{-1}$. There is an integer L depending only on ε such that $K^{2L} > |G|$. Then $\|\mu^{*(2^L l_1)}\|_2^2 = a_L \leq 2(\dim \pi)^{2/3} |G|^{-1}$.

Set $\mu_1 = \mu^{*(2^L l_1)}$, and consider the operator $T : f \mapsto \mu_1^{*(2)} * f$ acting on $L^2(G)$. We compute the trace of T . Recall that δ_g for $g \in G$ is the Dirac measure supported at g , and this constitute an orthonormal basis in $L^2(G)$. Hence

$$\begin{aligned} \text{Tr } T &= \sum_{g \in G} \langle T \delta_g, \delta_g \rangle = |G| \cdot \mu_1^{*(2)}(1) \\ &= |G| \sum_{g \in G} \mu_1(g) \mu_1(g^{-1}) = |G| \|\mu_1\|_2^2 \leq 2(\dim \pi)^{2/3}. \end{aligned}$$

We can also write $\text{Tr } T = \lambda_1 + \dots + \lambda_{|G|}$ as the sum of the eigenvalues of T . We can decompose the space $L^2(G)$ into the orthogonal sum of irreducible G representations. The number of components isomorphic to π in this decomposition is $\dim \pi$. Hence all eigenvalues of $\pi(\mu_1^{*(2)})$ occur with multiplicity at least $\dim \pi$ among the eigenvalues of T . Thus

$$\|\pi(\mu)\|^{2^{L+1}l_1} = \|\pi(\mu_1^{*(2)})\| \leq \frac{2(\dim \pi)^{2/3}}{\dim \pi}.$$

Taking this inequality to the $1/2^{L+1}l_1$ power, we get the conclusion of the theorem.

REFERENCES

- [BFGM07] U. Bader, A. Furman, T. Gelander, and N. Monod, *Property (T) and rigidity for actions on Banach spaces*, Acta Math. **198** (2007), no. 1, 57–105. MR2316269 (2008g:22007)
- [BG08] J. Bourgain and A. Gamburd, *Uniform expansion bounds for Cayley graphs of $SL_2(\mathbb{F}_p)$* , Ann. of Math. (2) **167** (2008), no. 2, 625–642. MR2415383 (2010b:20070)
- [BG09] J. Bourgain and A. Gamburd, *Expansion and random walks in $SL_d(\mathbb{Z}/p^n\mathbb{Z})$. II*, J. Eur. Math. Soc. (JEMS) **11** (2009), no. 5, 1057–1103. With an appendix by Bourgain. MR2538500 (2011a:60021)
- [BG10] E. Breuillard and A. Gamburd, *Strong uniform expansion in $SL(2, p)$* , Geom. Funct. Anal. **20** (2010), no. 5, 1201–1209. MR2746951 (2011m:20111)
- [BGT11] E. Breuillard, B. Green, and T. Tao, *Approximate subgroups of linear groups*, Geom. Funct. Anal. **21** (2011), no. 4, 774–819. MR2827010
- [BL00] Y. Benyamini and J. Lindenstrauss, *Geometric nonlinear functional analysis. Vol. 1*, American Mathematical Society Colloquium Publications, vol. 48, American Mathematical Society, Providence, RI, 2000. MR1727673 (2001b:46001)
- [Gow08] W. T. Gowers, *Quasirandom groups*, Combin. Probab. Comput. **17** (2008), no. 3, 363–387. MR2410393 (2009f:20105)
- [Hel08] H. A. Helfgott, *Growth and generation in $SL_2(\mathbb{Z}/p\mathbb{Z})$* , Ann. of Math. (2) **167** (2008), no. 2, 601–623. MR2415382 (2009i:20094)
- [Hel11] H. A. Helfgott, *Growth in $SL_3(\mathbb{Z}/p\mathbb{Z})$* , J. Eur. Math. Soc. (JEMS) **13** (2011), no. 3, 761–851. MR2781932
- [Kow12] E. Kowalski, *Crible en expansion*, Astérisque **348** (2012).
- [Kow13] E. Kowalski, *Explicit growth and expansion for SL_2* , Int. Math. Res. Not. IMRN **24** (2013), 5645–5708. MR3144176
- [LS74] V. Landazuri and G. M. Seitz, *On the minimal degrees of projective representations of the finite Chevalley groups*, J. Algebra **32** (1974), 418–443. MR0360852 (50 #13299)
- [Lub12] A. Lubotzky, *Expander graphs in pure and applied mathematics*, Bull. Amer. Math. Soc. (N.S.) **49** (2012), no. 1, 113–162. MR2869010 (2012m:05003)
- [LV14a] E. Lindenstrauss and P. P. Varju, *Spectral gap in the group of affine transformations over prime fields*, arXiv preprint arXiv:1409.3564v1 (2014). 26pp.
- [LV14b] E. Lindenstrauss and P. P. Varjú, *Spectral gap in the group of Euclidean isometries*, 2014.
- [LV] E. Lindenstrauss and P. P. Varjú, *Lectures on dynamical aspects of arithmetic combinatorics*. Work in progress.
- [LW93] A. Lubotzky and B. Weiss, *Groups and expanders*, Expanding graphs (Princeton, NJ, 1992), 1993, pp. 95–109. MR1235570 (95b:05097)
- [NP11] N. Nikolov and L. Pyber, *Product decompositions of quasirandom groups and a Jordan type theorem*, J. Eur. Math. Soc. (JEMS) **13** (2011), no. 4, 1063–1077. MR2800484 (2012h:20054)
- [PS10] L. Pyber and E. Szabó, *Growth in finite simple groups of Lie type of bounded rank*, 2010.

- [SGV12] A. Salehi Golsefidy and P. P. Varjú, *Expansion in perfect groups*, Geom. Funct. Anal. **22** (2012), no. 6, 1832–1891.
- [SX91] P. Sarnak and X. X. Xue, *Bounds for multiplicities of automorphic representations*, Duke Math. J. **64** (1991), no. 1, 207–227. MR1131400 (92h:22026)
- [Tao08] T. Tao, *Product set estimates for non-commutative groups*, Combinatorica **28** (2008), no. 5, 547–594. MR2501249 (2010b:11017)
- [Var12] P. P. Varjú, *Expansion in $SL_d(\mathcal{O}_K/I)$, I square-free*, J. Eur. Math. Soc. (JEMS) **14** (2012), no. 1, 273–305. MR2862040
- [Var] P. P. Varjú, *Random walks in Euclidean space*. Ann. of Math. (2), to appear.

(Elon Lindenstrauss) THE EINSTEIN INSTITUTE OF MATHEMATICS, EDMOND J. SAFRA CAMPUS, GIVAT RAM, THE HEBREW UNIVERSITY OF JERUSALEM, JERUSALEM, 91904, ISRAEL

E-mail address: elon@math.huji.ac.il

(Péter P. Varjú) UNIVERSITY OF CAMBRIDGE, DPMMS, WILBERFORCE ROAD, CAMBRIDGE, CB3 0WA, UK

E-mail address: pv270@dpmms.cam.ac.uk